arXiv:1611.02528v1 [cs.LO] 11 Oct 2016

LTL Model-Checking for Dynamic Pushdown Networks Communicating via Locks

Fu Song[1] and Tayssir Touili[2]

[1] Shanghai Key Laboratory of Trustworthy Computing, East China Normal University, Shanghai, China
[2] LIPN, CNRS and University Paris 13, France

Abstract. A Dynamic Pushdown Network (DPN) is a set of pushdown systems (PDSs) where each process can dynamically create new instances of PDSs. DPNs are a natural model of multi-threaded programs with (possibly recursive) procedure calls and thread creation. Extending DPNs with locks allows processes to synchronize with each other. Thus, DPNs with locks are a well adapted formalism to model multi-threaded programs that synchronize via locks. Therefore, it is important to have model-checking algorithms for DPNs with locks. We consider in this work model-checking for DPNs with locks against single-indexed LTL properties of the form $\bigwedge f_i$ s.t. f_i is a LTL formula interpreted over the PDS i. We consider the model-checking problems w.r.t. simple valuations (i.e, whether a configuration satisfies an atomic proposition depends only on its control location and held locks) and w.r.t. regular valuations (i.e., the set of the configurations satisfying an atomic proposition is a regular set of configurations). We show that these model-checking problems are decidable.

1 Introduction

It is notoriously difficult to write multithreaded programs whose bugs are commonly concurrency-related and are hard to reproduce and fix. Dynamic Pushdown Networks (DPN) [3] are a natural model of multi-threaded programs with (possibly recursive) procedure calls and thread creation. A DPN consists of a finite set of pushdown systems (PDSs), each of them models a sequential program (process) that can dynamically create new instances of PDSs. The model-checking programs (LTL, CTL or reachability properties) of DPNs are well studied in the literature [3,20,18,8,16,25,22]. DPNs with locks is an extension of DPNs in which processes can synchronize with each other via locks. DPNs with locks allow to model multi-threaded programs communicating via locks. However, only reachability properties are studied for DPNs with locks [18,8]. In this work, we consider model checking Linear Temporal Logic (LTL) which can describe more interesting properties of program behaviors.

In general, model checking DPNs for double-indexed properties (i.e., properties where atomic propositions are interpreted over the control states of two or more threads) or model checking DPNs with locks for reachabillity properties is undecidable [12]. This undecidability holds for pushdown networks even without thread creation. To obtain decidable results, in this paper, we consider single-indexed LTL properties for DPNs using locks in a nested style (L-DPN), where a single-index LTL formula is a

formula of the form $\bigwedge f_i$ such that f_i is a LTL formula interpreted over the PDS i, and using locks in a nested style allows each process to release only the latest acquired lock that has not yet been released. A L-DPN satisfies $\bigwedge f_i$ iff every PDS i running in parallel in the network satisfies the subformula f_i and the lock criteria [18].

We consider single-indexed LTL model-checking for L-DPNs with simple valuations (where whether a configuration of a PDS i satisfies an atomic proposition depends only on the control location of the configuration and its held locks) and with regular valuations (where the set of configurations of a PDS satisfying an atomic proposition is a regular set of configurations). We show that these model-checking problems are decidable. It is non-trivial to do LTL model checking for L-DPNs, since the number of instances of PDSs can be unbounded. Checking independently whether all the different PDSs satisfy the corresponding subformula f_i is not correct. Indeed, we do not need to check whether an instance of a PDS j satisfies f_j if this instance is not created during a run, and we have to guarantee that all the created instances use locks in a nested style. In our previous work [22], we have shown how to solve single-indexed LTL model-checking for L-DPNs without locks, i.e., DPNs. However, the approach of [22] cannot be directly applied to perform single-indexed LTL model-checking for L-DPNs due to locks. Indeed, we have to consider communication between each instance running in parallel in the network. To overcome this problem, inspired by the work of [19] which reduces reachability checking of L-DPNs to checking reachability of DPNS, we reduce single-indexed LTL model-checking for L-DPNs to single-indexed LTL model-checking for DPNs. For this, we will compute a DPN, a kind of "product" of the L-DPN with acquisition structures, where each control location of the DPN stores an acquisition structure. The acquisition structures stored in the control locations allow us to infer whether the global run uses locks correctly. For this, we characterize the set of consistent acquisition structures that will not violate the lock usages. The transformation of the acquisition structure during the global run of the DPN checks whether the lock usages is violated or not (i.e., the acquisition structure is consistent) and the runs of the DPN mimics the global run of the L-DPN. We disallow the global runs of the DPN in which an inconsistent acquisition structure will meet. By doing this, the global runs of the DPN exactly correspond to the global runs of the L-DPN that uses locks correctly. Thus, we reduce single-indexed LTL model-checking for L-DPNs to single-indexed LTL model-checking for DPNs. This later problem can be solved by our previous work [22].

Outline. Section 2 gives the basic definitions, recalls the results of single-indexed LTL model-checking for DPNs and presents a motivating example. Section 3 shows how to solve single-indexed LTL model-checking for L-DPNs. Section 4 discusses related works.

2 Preliminaries

2.1 LTL and Büchi Automata

From now on, we fix a set of atomic propositions AP.

Definition 1. *The set of LTL formulas is given by (where* $ap \in \mathrm{AP}$*):*

$$\psi \ ::= ap \mid \neg\psi \mid \psi \wedge \psi \mid \boldsymbol{X}\psi \mid \psi \boldsymbol{U} \psi.$$

Given an ω-word $\eta = \alpha_0\alpha_1...$ over 2^{AP}, let $\eta(k)$ denote α_k, and η_k denote the *suffix* of η starting from α_k. $\eta \models \psi$ (η satisfies ψ) is inductively defined as follows: $\eta \models ap$ iff $ap \in \eta(0)$; $\eta \models \neg\psi$ iff $\eta \not\models \psi$; $\eta \models \psi_1 \wedge \psi_2$ iff $\eta \models \psi_1$ and $\eta \models \psi_2$; $\eta \models \mathbf{X}\psi$ iff $\eta_1 \models \psi$; $\eta \models \psi_1 \mathbf{U} \psi_2$ iff there exists $k \geq 0$ such that $\eta_k \models \psi_2$ and for every j, $1 \leq j < k$, $\eta_j \models \psi_1$.

Definition 2. *A* Büchi automaton *(BA)* $\mathcal{B}$ *is a tuple* $(G, \Sigma, \theta, g^0, F)$ *where* G *is a finite set of states,* Σ *is the input alphabet,* $\theta \subseteq G \times \Sigma \times G$ *is a finite set of transitions,* $g^0 \in G$ *is the initial state and* $F \subseteq G$ *is a finite set of accepting states.*

A run of $\mathcal{B}$ over an ω-word $\alpha_0\alpha_1...$ is a sequence of states $q_0q_1...$ s.t. $q_0 = g^0$ and $(q_i, \alpha_i, q_{i+1}) \in \theta$ for every $i \geq 0$. A run is accepting iff it infinitely often visits some states in F.

It is well-known that given a LTL formula f, one can construct a BA B_f s.t. $\Sigma = 2^{AP}$ recognizing all the ω-words that satisfy f [24].

2.2 Dynamic Pushdown Networks with Locks

Definition 3. *A* Dynamic Pushdown Network *with Locks (L-DPN)* $\mathcal{M}$ *is a tuple* $(Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n)$ *s.t.* $\mathbb{L}$ *is a finite set of locks, Act is a finite set of actions* $\{acq(l), rel(l), \tau \mid l \in \mathbb{L}\}$ *where the action* $acq(l)$ *(resp.* $rel(l)$*) for every* $l \in \mathbb{L}$ *denotes the* acquisition *(resp.* release*) of the lock* l *and the action* τ *denotes all the lock-unrelated* internal actions*; for every* i, $1 \leq i \leq n$, $\mathcal{P}_i = (P_i, \Gamma_i, \Delta_i)$ *is a* Dynamic Pushdown System *(DPDS), where* P_i *is a finite set of control states s.t.* $P_k \cap P_i = \emptyset$ *for* $k \neq i$, Γ_i *is the stack alphabet,* Δ_i *is a finite set of transition rules in the following forms: (I)* $p_0\gamma \overset{a}{\hookrightarrow}_i p_1\omega_1$ *or (II)* $p_0\gamma \overset{a}{\hookrightarrow}_i p_1\omega_1 \rhd p_2\omega_2$ *s.t.* $a \in Act, p_0, p_1 \in P_i, \gamma \in \Gamma_i, \omega_1 \in \Gamma_i^*, p_2\omega_2 \in P_j \times \Gamma_j^*$ *for some* j, $1 \leq j \leq n$.

A L-DPN $\mathcal{M}$ *is a Dynamic Pushdown Network (DPN) if* $Act = \{\tau\}$ *and* $\mathbb{L} = \emptyset$*. We will write a DPN as* $(\mathcal{P}_1, ..., \mathcal{P}_n)$ *and sometimes we omit the labeling* τ *from all the transition rules in DPNs.*

For every $i : 1 \leq i \leq n$, a *local configuration* of a DPDS $\mathcal{P}_i$ is a tuple $(p\omega, L)$ such that $L \subseteq \mathbb{L}$ is a set of held locks, $p \in P_i$ is the control location and $\omega \in \Gamma_i^*$ is the stack content. Note that, if $Act = \{\tau\}$ and $\mathbb{L} = \emptyset$, a local configuration is denoted by $p\omega$. A DPDS is a pushdown system if all the transition rules are in the form of $q\gamma \overset{a}{\hookrightarrow}_i p_1\omega_1$.

A *global configuration* of $\mathcal{M}$ is a multiset over $\bigcup_{i=1}^{n} P_i \times \Gamma_i^* \times 2^{\mathbb{L}}$, in which each element $(p\omega, L)$ denotes the local configuration of an instance running in parallel in the network. Given a global configuration $\mathcal{G}$, the set of held locks at the global configuration is $(\bigcup_{(p\omega,L)\in\mathcal{G}} L)$, denoted by $hold(\mathcal{G})$. The set of free locks at the global configuration $\mathcal{G}$ is $\mathbb{L} \setminus hold(\mathcal{G})$, denoted by $free(\mathcal{G})$. Let $C_{\mathcal{M}}$ denote the set of global configurations of $\mathcal{M}$.

Given a tree T over $(\bigcup_{i=1}^{n} P_i \times \Gamma_i^*) \times 2^{\mathbb{L}}$, a node in the tree T that does not have any child is a leaf. Let $leaves(T)$ be the multiset that contains exactly the leaves of the tree T. W.l.o.g., we assume that the initial global configuration $\mathcal{G}$ of $\mathcal{M}$ contains only one

element. A *global run* of $\mathcal{M}$ starting from an initial global configuration $\mathcal{G}$ containing the element $(p_0\omega_0, L_0)$ is a binary tree T rooted by $(p_0\omega_0, L_0)$ and the leaves $leaves(T)$ of T is the current global configuration. The progress of the global run T is defined as follows: for every local configuration $(p\gamma u, L) \in leaves(T)$ of an instance of a DPDS $\mathcal{P}_i$ running in parallel in the network for some $\gamma \in \Gamma_i$:

α_1 : If there exists a transition rule $p\gamma \overset{\tau}{\hookrightarrow}_i p'\omega \in \Delta_i$, then $(p'\omega u, L)$ can be the right child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L)$, replacing the control location p by p' and the stack content γu by ωu without changing the set of held locks L. During this step, the other local configurations $leaves(T) \setminus (p\gamma u, L)$ of other instances running in parallel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L)\} \cup (leaves(T) \setminus (p\gamma u, L))$.

α_2 : If there exists a transition rule $p\gamma \overset{\tau}{\hookrightarrow}_i p'\omega \triangleright p_2\omega_2 \in \Delta_i$ s.t. $p_2 \in P_j$, then $(p'\omega u, L)$ and $(p_2\omega_2, \emptyset)$ can be the right child and left child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L)$. Moreover, a new instance of the DPDS $\mathcal{P}_j$ is created and it starts from the local configuration $(p_2\omega_2, \emptyset)$. Here, we suppose w.l.o.g., that the set of locks held by this new instance is empty. This kind of newly created instance during the global run is *called dynamically created new instance*, and a local configuration that a dynamically created new instance starts from is called by *Dynamically Created Local Initial Configuration* (DCLIC for short). The other local configurations $leaves(T)\setminus(p\gamma u, L)$ of other instances running in parallel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L), (p_2\omega_2, \emptyset)\} \cup (leaves(T) \setminus (p\gamma u, L))$.

α_3 : If there exists a transition rule $p\gamma \overset{acq(l)}{\hookrightarrow}_i p'\omega \in \Delta_i$ such that $l \in free(leaves(T))$ (i.e., l is a free lock at the current global configuration $leaves(T)$), then $(p'\omega u, L \cup \{l\})$ can be the right child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L \cup \{l\})$ and hold the lock l. During this step, the other local configurations $leaves(T) \setminus (p\gamma u, L)$ of other instances running in parallel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L \cup \{l\})\} \cup (leaves(T) \setminus (p\gamma u, L))$.

α_4 : If there exists a transition rule $p\gamma \overset{acq(l)}{\hookrightarrow}_i p'\omega \triangleright p_2\omega_2 \in \Delta_i$ such that $l \in free(leaves(T))$, then $(p'\omega u, L \cup \{l\})$ and $(p_2\omega_2, \emptyset)$ can be the right child and left child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L \cup \{l\})$, hold the lock l and create a new instance starting from $(p_2\omega_2, \emptyset)$. During this step, the other local configurations $leaves(T) \setminus (p\gamma u, L)$ of other instances running in parallel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L \cup \{l\}), (p_2\omega_2, \emptyset)\} \cup (leaves(T) \setminus (p\gamma u, L))$.

α_5 : If there exists a transition rule $p\gamma \overset{rel(l)}{\hookrightarrow}_i p'\omega \in \Delta_i$ such that $l \in L$ (i.e., the instance owns the lock l), then $(p'\omega u, L \setminus \{l\})$ can be the right child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L \setminus \{l\})$ and free the lock l. The other local configurations $leaves(T)\setminus(p\gamma u, L)$ of other instances running in par-

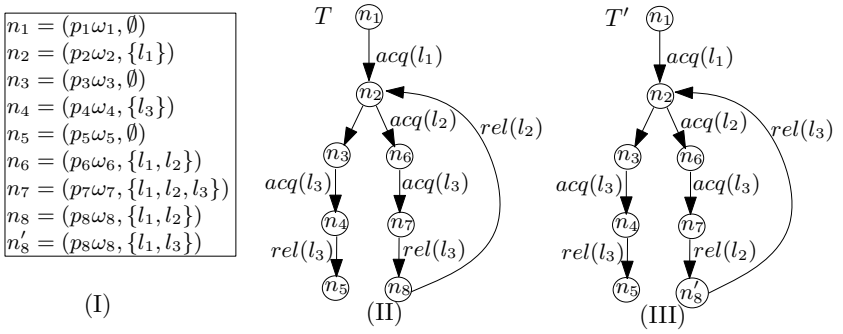

Fig. 1. (a) and (b) are two global runs that using locks (not) in a nested style.

allel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L \setminus \{l\})\} \cup (leaves(T) \setminus (p\gamma u, L))$.

α_6 : If there exists a transition rule $p\gamma \overset{rel(l)}{\hookrightarrow}_i p'\omega \triangleright p_2\omega_2 \in \Delta_i$ such that $l \in L$ (i.e., the instance owns the lock l), then $(p'\omega u, L \setminus \{l\})$ can be the right child of the node $(p\gamma u, L)$. This means that this instance can move from $(p\gamma u, L)$ to $(p'\omega u, L \setminus \{l\})$, free the lock l and create a new instance starting from $(p_2\omega_2, \emptyset)$. The other local configurations $leaves(T) \setminus (p\gamma u, L)$ of other instances running in parallel in the network stay at the same local configurations. After applying this transition rule, the next global configuration is $\{(p'\omega u, L \setminus \{l\}), (p_2\omega_2, \emptyset)\} \cup (leaves(T) \setminus (p\gamma u, L))$.

Intuitively, each left child in a global run T is a local initial configuration of the newly created instance and the root is the local initial configuration of the initial instance (i.e., the instance is not created on the runtime). Each rightmost path is a trance of an instance running in parallel in the global run (network). The rightmost path starting from the root or a left child is a *local run* of the initial instance or a newly created instance running in parallel in the network, respectively. Note the defining the global runs as trees allows us to know which instance creates a new instance, where it creates and the local initial configuration of the newly created instance. It is important to reasoning about lock usages. Let $\mathcal{D}_I = \{(p_2\omega_2, \emptyset) \mid p\gamma \overset{a}{\hookrightarrow}_i p'\omega \triangleright p_2\omega_2 \in \Delta_i \text{ for some } i, 1 \leq i \leq n\}$ be the set of all the possbile DCLICs. When the L-DPN is DPN, $\mathcal{D}_I = \{p_2\omega_2 \mid p\gamma \overset{a}{\hookrightarrow}_i p'\omega \triangleright p_2\omega_2 \in \Delta_i \text{ for some } i, 1 \leq i \leq n\}$.

Nested Lock Access. A global run of $\mathcal{M}$ uses locks in a nested style iff each local run running in parallel in the global run uses locks in a nested style, i.e., the local run releases only the latest acquired lock that has not yet been released. In this work, we consider L-DPNs that use locks in a nested style. This is because even reachability, and hence LTL, is known to be undecidable for pushdown networks that use locks in an arbitary style [12].

Example 1. Figure 1(II) and Figure 1(III) show two global runs called by T and T' that uses locks l_1, l_2 and l_2. Each edge is labeled by the corresponding action. The nodes $n_1, ..., n_8$ and n'_8 denote the local configurations showing in Figure 1(I). $n_1(n_2n_6n_7n_8)^*$

(resp. $n_1(n_2n_6n_7n_8')^*$) is a local run in T (resp. T') during which a new instance is created when moving from n_2 to n_6. This new instance has the local run $n_3n_4n_5$. In T, we can see that locks are accessed in a nested style. While, in T', the locks l_1 and l_2 are not accessed in a nested style, since the lock l_2 is released before the release of the latest acquired lock l_3.

W.l.o.g., in this work, we consider only non-reentrant locks, i.e., a process cannot acquire the same lock multiple times before releasing it. While reentrant locks allow a process to acquire the same lock multiple times before releasing it. Indeed, reentrant locks can be simulated with non-reentrant locks [18].

2.3 Representing Infinite Set of Local Configurations

To finitely represent (infinite) sets of local configurations of DPDSs and DCLICs generated by DPDSs, we introduce L-Multi-automata and Multi-automata.

Definition 4. *Given a L-DPN $\mathcal{M} = (Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n)$, a L-Multi-automaton* (L-MA) *is a tuple $\mathcal{A}_i = (Q_i, \Gamma_i, \delta_i, I_i, Acc_i)$, where Q_i is a finite set of states, $I_i \subseteq P_i \times 2^{\mathbb{L}}$ is a finite set of initial states corresponding to the control locations and held locks of the DPDS $\mathcal{P}_i$, $Acc_i \subseteq Q_i$ is a finite set of final states, $\delta_i \subseteq (Q_i \times \Gamma_i) \times 2^{\mathcal{D}_I} \times Q_i$ is a finite set of transition rules. A L-MA is a Multi-automaton (MA) if the L-DPN is a DPN, i.e., $I_i \subseteq P_i$ (note that $\mathbb{L} = \emptyset$).*

We write $q \xrightarrow{\gamma/D}_i q'$ instead of $(q, \gamma, D, q') \in \delta_i$, where D is a set of DCLICs. We define the relation $\longrightarrow_i^* \subseteq (Q_i \times \Gamma_i^*) \times 2^{\mathcal{D}_I} \times Q_i$ as the smallest relation s.t.: (1) $q \xrightarrow{\epsilon/\emptyset}_i^* q$ for every $q \in Q_i$, (2) if $q \xrightarrow{\gamma/D_1}_i q_1$ and $q_1 \xrightarrow{\omega/D_2}_i^* q_2$, then $q \xrightarrow{\gamma\omega/D_1 \cup D_2}_i^* q_2$. Let $L(\mathcal{A}_i)$ be the set of tuples $(p\omega, L, D) \in P_i \times \Gamma_i^* \times 2^{\mathcal{D}_I}$ s.t. $(p, L) \xrightarrow{\omega/D}_i^* q$ for some $q \in Acc_i$. A set $W \subseteq P_i \times \Gamma_i^* \times 2^{\mathbb{L}} \times 2^{\mathcal{D}_I}$ is *regular* iff there exists a L-MA $\mathcal{A}_i$ s.t. $L(\mathcal{A}_i) = W$. A set of local configurations $C \subseteq P_i \times \Gamma_i^* \times 2^{\mathbb{L}}$ is *regular* iff $C \times \{\emptyset\}$ is a regular set.

2.4 Single-indexed LTL Model-Checking for L-DPNs

Model-checking L-DPNs for double-indexed LTL properties where the validity of atomic propositions depends on two or more DPDSs is undecidable. Indeed, model-checking double-indexed LTL properties for pushdown networks even without interaction with each other is undecidable [12]. Thus, in this work, we consider model-checking L-DPNs for single-indexed LTL properties of the form $\bigwedge_{i=1}^n f_i$ s.t. for every $i : 1 \leq i \leq n$, f_i is a LTL formula interpreted over the DPDS $\mathcal{P}_i$. From now on, we fix a L-DPN $\mathcal{M} = (Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n$ s.t. for every i, $1 \leq i \leq n$, $\mathcal{P}_i = (P_i, \Gamma_i, \Delta_i)$ and a single-index LTL formula $\bigwedge_{i=1}^n f_i$ such that for every $i : 1 \leq i \leq n$, f_i is interpreted over the DPDS $\mathcal{P}_i$.

Given a valuation $\lambda : AP \longrightarrow 2^{\bigcup_{i=1}^n (P_i \times \Gamma_i^* \times 2^{\mathbb{L}})}$ that assigns to each atomic proposition a set of local configurations, and a global run T of the L-DPN $\mathcal{M}$, a local run $\sigma = (p_0\omega_0, L_0)(p_1\omega_1, L_1)...$ of an instance of $\mathcal{P}_i$ running in parallel in the global run T satisfies f_i iff the ω-word $\sigma_0\sigma_1...$ satisfies f_i, where for every $j \geq 0$,

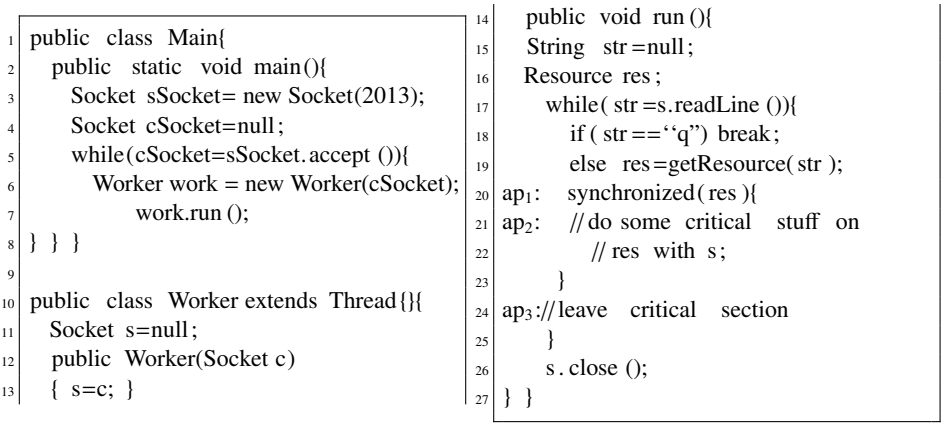
1 public class Main{
2 public static void main(){
3 Socket sSocket= new Socket(2013);
4 Socket cSocket=null;
5 while(cSocket=sSocket.accept ()){
6 Worker work = new Worker(cSocket);
7 work.run ();
8 } } }
9
10 public class Worker extends Thread{}{
11 Socket s=null;
12 public Worker(Socket c)
13 { s=c; }
14 public void run (){
15 String str=null;
16 Resource res;
17 while(str=s.readLine ()){
18 if (str==“q”) break;
19 else res=getResource(str);
20 ap1: synchronized(res){
21 ap2: // do some critical stuff on
22 // res with s;
23 }
24 ap3://leave critical section
25 }
26 s.close ();
27 } }

Fig. 2. A Java-like concurrent server program. a, b and c atomic propositions associated with the corresponding control points.

$\sigma_j = \{ap \in AP \mid (p_j\omega_j, L_j) \in \lambda(at)$. A global run T of the L-DPN $\mathcal{M}$ satisfies f iff each local run of each instance of $\mathcal{P}_i$ for $1 \leq i \leq n$ running in parallel in the network satisfies f_i. A global initial configuration $\mathcal{G}$ satisfies f iff $\mathcal{M}$ has a global run T starting from $\mathcal{G}$ such that T satisfies f.

However, checking whether a local run of an instance of a pushdown system $\mathcal{P}_i$ without locks satisfies the formula f_i with respect to an arbitary valuation λ is undecidable [7]. To have decidable results, in this work, we first consider L-DPNs model-checking for single-indexed LTL with simple valuations, where the validity of each atomic proposition only depends on the control location and locks. Formally, a simple valuation is a function $\lambda : AP \longrightarrow 2^{\cup_{i=1}^{n}(P_i \times 2^{\mathbb{L}})}$ that assigns to each atomic proposition a set of control locations and locks. It can be extended to $\lambda : AP \longrightarrow 2^{\cup_{i=1}^{n}(P_i \times \Gamma_i^* \times 2^{\mathbb{L}})}$ as usual.

Next, we consider a more general problem, single-indexed LTL model-checking problem for L-DPNs with regular valuations, where the validity of each atomic proposition depends on the control locations, the set of locks and the stack content. Since a local configuration consists of a control location, the stack content and a set of held locks, it is natural to consider this model-checking problem. Formally, a regular valuation is a function $\lambda : AP \longrightarrow 2^{\cup_{i=1}^{n}(P_i \times \Gamma_i^* \times 2^{\mathbb{L}})}$ such that for every $ap \in AP$, the set of local configurations that satisfies ap is a regular set of local configurations, i.e., there is a L-MA M such that $(p\omega, L) \in \lambda(ap)$ iff $(p\omega, L, \emptyset) \in L(M_{ap})$.

2.5 A Motivating Example

Figure 2 shows a simplified Java-like concurrent server program. The main process creates a socket object *sSocket* to listen on the port *2013*. When a request (connection) arrives from a client, the main process creates a new process *work*, passes the connection *cSocket* to *work*, executes the new process *work* by invoking *run()* and continues

listening on the port *2013*. The process *work* accepts a request from the main process and processes it. If it receives "q" from the client, it closes the connection by invoking *close()* and this process terminates. Otherwise, it obtains the resource *res* according to the request *str* and checks for availability of the lock l implicitly associated with *res* by *synchronized(res)*. If l is free, *work* successfully acquires l, performs the critical operations, after that l is released. Otherwise, if l has already been acquired by another process *work'*, then *work* becomes blocked until l is free. This program has an unbound number of processes of *Worker* which communicate with other processes via nested locks.

L-DPNs are well-suitable to model such kind of concurrent programs. Several interesting properties of this program can be expressed in single-indexed LTL formulas. Suppose the scheduler is fair. ap_1, ap_2 and ap_3 atomic propositions associated with the corresponding control points. The *starvation* property can be expressed as $\mathbf{F}ap_1 \wedge \mathbf{G}\neg ap_2$ checking whether there is a local run of *Worker* running in parallel in the network such that at some point, ap_1 holds but ap_2 will never hold in the future, i.e., the local run will never enter the critical section. The *deadlock* property can be expressed as $\mathbf{FG}ap_1$ checking whether or not the lock l can be held forever by some process. The property that each process always leave the synchronized block once it enters can be expressed as $\mathbf{F}ap_2 \wedge \mathbf{G}\neg ap_3$. The *mutual exclusion* over the shared resource can also be checked by our techniques. We can create two DPDSs $\mathcal{P}$ and $\mathcal{P}'$ for *Worker* in which the atomic proposition ap_2 is named by ap_2' in $\mathcal{P}'$. The main process alternatively creates new processes using $\mathcal{P}$ and $\mathcal{P}'$ when a request arriving. The mutual exclusion property is expressed as $\mathbf{F}ap_2 \wedge \mathbf{F}ap_2'$.

2.6 Single-indexed LTL Model-Checking for Dynamic Pushdown Networks

Single-index LTL model-checking for DPNs with simple and regular valuations was studied in [22]. We recall these results in this section, since we will reduce single-indexed LTL model-checking for L-DPNs to these problems. Let $\pi(p)$ denote the index i such that $p \in P_i$.

Theorem 1. *(Thm. 3 of [22]) We can compute a set $\mathcal{D}_{fp}$ of DCLICs in time $\mathbf{O}(\sum_{p\omega\in\mathcal{D}_I}(|\omega| \cdot |\delta_{\pi(p)}| \cdot |Q_{\pi(p)}|) \cdot 2^{|\mathcal{D}_I|} + |\mathcal{D}_I|^2 \cdot 2^{|\mathcal{D}_I|})$ s.t. for every $c \in \mathcal{D}_I$, c satisfies the single-indexed LTL formula f iff $c \in \mathcal{D}_{fp}$.*

Theorem 2. *(Thm. 4 of [22]) Given a DPN $\mathcal{M} = \{\mathcal{P}_1, ..., \mathcal{P}_n\}$, a single-indexed LTL formula $f = \bigwedge_{i=1}^{n} f_i$ and a simple function λ, we can compute MAs $\mathcal{A}_1, ..., \mathcal{A}_n$ in time $O(\sum_{i=1}^{n}(|\Delta_i| \cdot 2^{|f_i|} \cdot |\Gamma_i| \cdot |P_i|^3) \cdot 2^{|\mathcal{D}_I|})$ s.t. for every global configuration $p\omega$, $p\omega$ satisfies f iff there exists $D \subseteq \mathcal{D}_{fp}$ s.t. $(p\omega, D) \in L(\mathcal{A}_{\pi(p)})$.*

Theorem 3. *(Thm. 5 of [22]) Given a DPN $\mathcal{M} = \{\mathcal{P}_1, ..., \mathcal{P}_n\}$, a single-indexed LTL formula $f = \bigwedge_{i=1}^{n} f_i$ and a regular valuation λ, we can compute MAs $\mathcal{A}_1, ..., \mathcal{A}_n$ in time $\mathbf{O}(\sum_{i=1}^{n}(|\Delta_i| \cdot 2^{|f_i|} \cdot |\Gamma_i| \cdot |States_i| \cdot |P_i|^3) \cdot 2^{|\mathcal{D}_I|})$ s.t. for every global configuration $p\omega$, $p\omega$ satisfies f iff there exists $D \subseteq \mathcal{D}_{fp}$ s.t. $(p\omega, D) \in L(\mathcal{A}_{\pi(p)})$, where $|States_i|$ denotes the number of states of the automata corresponding to the regular valuation λ.*

3 Single-indexed LTL Model-Checking for L-DPNs

To check whether a L-DPN $\mathcal{M}$ satisfies f is non-trivial, we cannot directly apply the approach of [22] to check whether the L-DPN $\mathcal{M}$ satisfies f or not, as we have to ensure that the access of locks in each instance is correctly coordinated with other instances. To solve this problem, we follow the work of [19] which reduces reachability checking of L-DPNs to checking reachability of DPNS, We will reduce single-indexed LTL model-checking for L-DPNs to single-indexed LTL model-checking for DPNs. For this, for every $i, 1 \leq i \leq n$, we compute a new DPDS $\mathcal{P}'_i$, which is a kind of "product" of the DPDS $\mathcal{P}_i$ with acquisition structures, where an acquisition structure contains lock usage information such as the set of held locks, the order dependence of acquisition and release of locks. We will associate each control location of $\mathcal{P}'_i$ with an acquisition structure. The acquisition structure stored in a local configuration "guesses" the acquisition and release histories of locks in the global run starting from this local configuration. From the acquisition structures, we can infer the lock usages of a global run. Then, we can obtain a DPN $\mathcal{M}' = (\mathcal{P}'_1, ..., \mathcal{P}'_n)$ such that the global runs of $\mathcal{M}'$ mimic the global runs of $\mathcal{M}$. Since we can check lock usages from the acquisition structures, we disallow all the acquisition structures that violate the nested lock style in the global run of $\mathcal{M}'$. Thus, the global runs of $\mathcal{M}'$ exactly correspond to the global runs of $\mathcal{M}$ that use locks in a nested style. We can get that $\mathcal{M}$ satisfies f iff $\mathcal{M}'$ satisfies f. The later problem can be solved by Theorems 1, 2 and 3.

3.1 Acquisition Structures

Along a global run T of $\mathcal{M}$, a release of a lock l without a corresponding acquisition of l in the same local run is called *initial release*. An acquisition of a lock l without a corresponding release of l in the same local run is called *final acquisition*. An *acquisition structure as* of a global run T is a tuple (R, RH, U, AH, A, X), where R (resp. A) is the set of initial release (resp. final acquisition) locks of T; U is a set of *usages*, i.e., acquisition and release locks that are not final acquisition or initial release locks are called; $RH \subseteq \mathbb{L} \times \mathbb{L}$ is a *release graph* such that $(l, l') \in RH$ iff T has an initial release of l' and the usages of l occurs before the initial release of l'; $AH \subseteq \mathbb{L} \times \mathbb{L}$ is an *acquisition graph* such that $(l, l') \in AH$ iff T has a final acquisition of l and the usages of l' occurs after the final acquisition of l; X is a set of locks that are initially-held at the root.

An acquisition structure $as = (R, RH, U, AH, A, X)$ is *consistent* iff both RH and AH are acyclic, and $(X \setminus R) \cap (U \cup A) = \emptyset$. Intuitively, the set of locks $(X \setminus R)$ denotes all the initially-held locks that will not be released during the run. Thus, these locks cannot be used (i.e., $(X \setminus R) \cap U = \emptyset$) or finally acquired (i.e., $(X \setminus R) \cap A = \emptyset$) anymore. The fact that the graphs RH and AH are acyclic ensures that the acquisition and release of locks do not have any cycle dependence. Intuitively, if RH has edges $(l_1, l_2), ..., (l_m, l_{m+1})$ for some $m > 1$ such that $l_1 = l_m + 1$ (i.e., RH has a cycle), then for every $i : 1 < i \leq m + 1$, the lock l_i has an initial release in the global run T and should be performed after a usage of l_{i-1} (according to the definition of release graphs). Since the initial release of l_i releases the initially-held lock $l_i \in X$, then before releasing this initially-held lock l_i, l_i cannot be acquired anymore, i.e., the usage of l_i occurs after the initial release of l_i. Thus, the initial release of l_{i+1} should be done after the initial release

of l_i. Since $l_1 = l_{m+1}$, then, the initial release of l_1 should be performed after the initial release of l_1. We deduce a deadlock. Thus, RH should be acyclic. AH is similar. Let AS be the set of all the consistent acquisition structures. Given an acquisition structure $as = (R, RH, U, AH, A, X)$, we will write $as_R = R$, $as_{RH} = RH$, $as_U = U$, $as_{AH} = AH$, $as_A = A$ and $as_X = X$.

Example 2. Let us consider the global run T given in Figure 1(II). During the local run $n_1(n_2n_6n_7n_8)^*$, it will always create a new instance whose local run is $n_3n_4n_5$. Let (R, RH, U, AH, A, X) be the acquisition structure of the node n_1 in T. Suppose the set of initially-held locks in T is $\emptyset$, then $X = \emptyset$. Since all the releases in T has an corresponding acquisition, we can know that $R = \emptyset$ which implies that $RH = \emptyset$. While the acquisition of the locks except $acq(l_1)$ does not have any corresponding release, thus $A = \{l_1\}$. The acquisition of the locks l_2 and l_3 (i.e. $acq(l_2), acq(l_3)$) have corresponding releases (i.e., $rel(l_2), rel(l_3)$), we deduce that $U = \{b, c\}$. The order of the final acquisition of l_1 and the usages of l_2 and l_3 gives us that $AH = \{(l_1, l_2), (l_1, l_3)\}$. Similarly, the acquisition structure of T^{n_6} (i.e., the subtree rooted by the node n_6) is $(R', RH', U', AH', A', X')$, where $R' = \{l_2\}$ (since the initial release $rel(l_2)$ does not have any corresponding acquisition of l_2), $X' = \{l_1, l_2\}$ (the locks l_1 and l_2 acquired when moving from n_1 to n_6 are not yet released), $RH' = \{(l_3, l_2)\}$, $A' = AH' = \emptyset$, $U' = \{l_2, l_3\}$.

Given two consistent acquisition structures $as = (R_1, RH_1, U_1, AH_1, A_1, X_1)$ and $as' = (R_2, RH_2, U_2, AH_2, A_2, X_2)$, as and as' are *compatible*, denoted by $Compatible(as, as')$, iff the following conditions hold:

1. $X_1 \cap X_2 = \emptyset$;
2. $(A_1 \cup (X_1 \setminus R_1)) \cap (A_2 \cup (X_2 \setminus R_2)) = \emptyset$;
3. $RH_1 \cup RH_2$ is acyclic;
4. $AH_1 \cup AH_2$ is acyclic;
5. $(A_1 \cup U_1) \cap (X_2 \setminus R_2) = \emptyset$
6. $(A_2 \cup U_2) \cap (X_1 \setminus R_1) = \emptyset$.

Roughly specking, the compatible condition is used to check whether two global runs could be two subtrees of a node in a global run. Imagine there is a global run such that $(p\omega, L)$ is a leaf and $(p_1\omega_1, L_1)$ (resp. $(p_2\omega_2, L_2)$) could be the right (resp. left) child of the leaf $(p\omega, L)$ using a transition rule with τ action. Suppose T_1 and T_2 be two global runs that are rooted by $(p_1\omega_1, L_1)$ and $(p_2\omega_2, L_2)$, respectively. Let as^1 and as^2 be the acquisition structures of T_1 and T_2, respectively. The $compatible(as^1, as^2)$ checks whether the two tree T_1 and T_2 could be two subtrees of the leaf $(p\omega, L)$. Each condition verifies whether the nested lock access is violated or not. Items 1 and 2 verify that the initially-held locks and finally-held hocks are disjoint. If there is a lock $l \in X_1 \cap X_2$, then the lock l will be held by the two local configuration $(p_1\omega, L_1)$ and $(p_2\omega_2, L_2)$. This implies that T_1 and T_2 cannot be the subtrees of the node $(p\omega, L)$ in the global run of $\mathcal{M}$ due to a deadlock. On the other hand, if there is a lock $l \in (A_1 \cup (X_1 \setminus R_1)) \cap (A_2 \cup (X_2 \setminus R_2))$, then, T_1 (resp. T_2) has a final acquisition of the lock l or l is held at the initial global configuration of T_1 (resp. T_2) that will not be released. This means that both T_1 and T_2 will finally hold the lock l which is a deadlock. Items 3 and 4 verify that the acquisition and release graphs are acyclic. Since if there is a cycle in $RH_1 \cup RH_2$ or $AH_1 \cup AH_2$, as discussed previously, it will have a deadlock. Items 5 and 6 verify that the held throughout locks (the set of initially-held locks that will not be released) are not acquired or released. If there is a lock $l \in (A_1 \cup U_1) \cap (X_2 \setminus R_2)$, then, the lock l will always be held in T_2. This implies that T_1 should not acquire l_1 (i.e, $l_1 \notin A_1 \cup U_1$).

3.2 From L-DPN to DPN

In this section, we show how to compute a DPN $\mathcal{M}' = (\mathcal{P}'_1, ..., \mathcal{P}'_n)$, a kind of "product" of the L-DPN $\mathcal{M}$ with the acquisition structures AS such that a global run T' of $\mathcal{M}'$ mimics a corresponding global run T of $\mathcal{M}$, i.e., for every node $(p\omega, L)$ in T, there is a corresponding node $(p, as)\omega$ in T', where as is the acquisition structure of the subtree rooted by $(p\omega, L)$ in T. We update the acquisition structures embedded in the control locations of $\mathcal{P}'_i$ during the global run of $\mathcal{M}'$ and checks whether the acquisition structure is consistent or not. If an inconsistent acquisition structure occurs in a global run of $\mathcal{M}'$, then the corresponding global run of $\mathcal{M}$ violates the lock usages. Thus, we disallow all the global runs of $\mathcal{M}'$ in which an inconsistent acquisition structure occurs. Then, $\mathcal{M}$ has a global run T that starts from a global configuration $(p\omega, L)$, uses locks in a nested style and satisfies f iff $\mathcal{M}'$ has a corresponding global run T' starting from a global configuration $(p, as)\omega$ for some $as \in AS$ such that T' satisfies f.

To compute $\mathcal{M}'$, for every $i : 1 \leq i \leq n$, let $\mathcal{P}'_i = (P'_i, \Gamma_i, \Delta'_i)$, where $P'_i = P_i \times AS$ and Δ'_i is computed as follows: for every $as, as', as'' \in AS$,

1. $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \in \Delta'_i$ iff $p\gamma \overset{a}{\hookrightarrow}_i p'\omega \in \Delta_i$, and one of the following conditions holds:

   1.1: $as = as'$, if $a = \tau$; or

   1.2: $as = \Big(as'_R \cup \{l\}, as'_{RH}, as'_U, as'_{AH}, as'_A, as'_X \cup \{l\}\Big)$, if $a = rel(l)$ and $l \notin as'_X \cup as'_R$; or

   1.3: $as = \begin{cases} \Big(as'_R \setminus \{l\}, (as'_{RH} \setminus (\mathbb{L} \times \{l\}) \cup (\{l\} \times as'_R \setminus \{l\})), as'_U \cup \{l\}, as'_{AH}, as'_A, \\ as'_X \setminus \{l\}\Big), \qquad \text{if } a = acq(l) \text{ and } l \in as'_R \cap as'_X; \\ \\ \Big(as'_R, as'_{RH}, as'_U, as'_{AH} \cup \{\{l\} \times as'_U\}, as'_A \cup \{l\}, as'_X \setminus \{l\}\Big), \\ \qquad \text{else if } a = acq(l) \text{ and } l \notin as'_A \text{ and } l \in as'_X; \end{cases}$

2. $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \rhd (p_1, as'')\omega_1 \in \Delta'_i$ iff $p\gamma \overset{a}{\hookrightarrow}_i p'\omega \rhd p_1\omega_1 \in \Delta_i$, $Compatible(as', as'')$, $as''_R = as''_X = \emptyset$, and one of the following conditions holds:

   2.1: $as = (as'_R \cup as''_R, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, as'_X \cup as''_X)$, if $a = \tau$; or

   2.2: $as = \Big(as'_R \cup as''_R \cup \{l\}, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, as'_X \cup as''_X \cup \{l\}\Big)$, if $a = rel(l)$ and $l \notin as'_X \cup as''_X \cup as'_R \cup as''_R$; or

   2.3: $as = \begin{cases} \Big((as'_R \cup as''_R) \setminus \{l\}, ((as'_{RH} \cup as''_{RH}) \setminus (\mathbb{L} \times \{l\}) \cup (\{l\} \times (as'_R \cup as''_R) \setminus \{l\})), \\ as'_U \cup as''_U \cup \{l\}, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, (as'_X \cup as''_X) \setminus \{l\}\Big), \\ \qquad \text{if } a = acq(l) \text{ and } l \in (as'_R \cup as''_R) \cap (as'_X \cup as''_X); \\ \\ \Big(as'_R \cup as''_R, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, \{\{l\} \times (as'_U \cup as''_U)\} \cup as'_{AH} \cup as''_{AH}, \\ as'_A \cup as''_A \cup \{l\}, (as'_X \cup as''_X) \setminus \{l\}\Big), \\ \qquad \text{else if } a = acq(l) \text{ and } l \notin as'_A \cup as''_A \text{ and } l \in as'_X \cup as''_X. \end{cases}$

Intuitively, the DPN $\mathcal{M}' = (\mathcal{P}'_1, ..., \mathcal{P}'_n$ is a kind of "product" of $\mathcal{M}$ with acquisition structures AS such that $\mathcal{M}$ has a global run T starting from a global configuration $(p\omega, L)$ such that T uses locks in a nested style and satisfies f iff $\mathcal{M}'$ has a corresponding global run T' starting from a global configuration $(p, as)\omega$ for some $as \in AS$ such that T' satisfies f. Intuitively, if we reserve only the set of held locks X of each acquisition structure in the global runs of $\mathcal{M}'$, then the global runs of $\mathcal{M}'$ are exactly the global runs of $\mathcal{M}$. The acquisition structure updated during the global runs of $\mathcal{M}'$ ensures that the consistent acquisition structure as of the root $(p, as)\omega$ is the acquisition structure of the tree T, i.e., T uses locks in a nested style. Let us explain the intuition behind each items by structural induction on the global run of $\mathcal{M}$ ($\mathcal{M}'$). Suppose the global run T of $\mathcal{M}$ has a node $(p\omega, L) \in P_i \times \Gamma_i^* \times 2^{\mathbb{L}}$ s.t. $\omega = \gamma u$ for some $i : 1 \leq i \leq n$, $\gamma \in \Gamma_i$ and the subtree rooted by $(p\omega, L)$ is T_1.

If the node $(p\omega, L)$ has the right and left children $(p'vu, L)$ and $(p_2\omega_2, \emptyset)$ due to the transition rule $p\gamma \overset{\tau}{\hookrightarrow}_i p'v \triangleright p_2\omega_2 \in \Delta_i$, then we suppose that the subtree rooted by $(p'vu, L)$ (resp. $(p_2\omega_2, \emptyset)$) in T be T_2 (resp. T_3) and the acquisition structure of T_2 (resp. T_3) is as' (resp. as''). According to the definition of the acquisition structures, $as''_X = \emptyset$. This implies that $as''_R = \emptyset$, since there does not any lock need to release. If $Compatible(as', as'')$ (as explained in Section 3.1) does not hold or as' (resp. as'') is inconsistent, then the global run T violates the lock usages. This implies that all the global runs of $\mathcal{M}'$ should not contain a node $(p, as')'vu$ (resp. $(p_2, as'')\omega_2$)). Thus, we do not add a transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \triangleright (p_1, as'')\omega_1$ into Δ'_i for any $as \in AS$ if $Compatible(as', as'')$ does not hold, or as' (resp. as'') is inconsistent. Otherwise, according to the definition of acquisition structures, the acquisition structure as of T_1 must be $(as'_R \cup as''_R, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, as'_X \cup as''_X)$. This implies that if $(p, as)\gamma v$ is a node of a global run of $\mathcal{M}'$, then $(p', as')uv$ and $(p_1, as'')\omega_1$ can be the right and left children of $(p, as)\gamma v$. For this, we add the transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \triangleright (p_1, as'')\omega_1$ into Δ'_i.

If the node $(p\omega, L)$ has the right and left children $(p'vu, L \setminus \{l\})$ and $(p_2\omega_2, \emptyset)$ due to the transition rule $p\gamma \overset{rel(l)}{\hookrightarrow}_i p'v \triangleright p_2\omega_2 \in \Delta_i$, then $l \in L$ (see Item α_6). We suppose that the subtree rooted by $(p'vu, L)$ (resp. $(p_2\omega_2, \emptyset)$) in T be T_2 (resp. T_3) and the acquisition structure of T_2 (resp. T_3) is as' (resp. as''), then $l \notin as'_X$, $as''_X = \emptyset$, $as''_X = \emptyset$ and $as''_R = \emptyset$ according to the definition of the acquisition structures. This implies that T_2 and T_3 should not have any initial release of l, (i.e., $l \notin as'_R \cup as''_R$). Thus, if $l \in as'_X \cup as''_X \cup as'_R \cup as''_R$, we do not add a transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \triangleright (p_1, as'')\omega_1$ into Δ'_i for any $as \in AS$. Otherwise, if $Compatible(as', as'')$ does not hold or as' (resp. as'') is inconsistent, then we neither do not add a transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \triangleright (p_1, as'')\omega_1$ into Δ'_i for any $as \in AS$ as explained above. If $l \notin as'_X \cup as''_X \cup as'_R \cup as''_R$ and $Compatible(as', as'')$ holds, according to the definition of acquisition structures, the acquisition structure as of T_1 must be $\Big(as'_R \cup as''_R \cup \{l\}, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, as'_X \cup as''_X \cup \{l\}\Big)$ (note that $rel(l)$ is an initial release for T_1). This implies that if $(p, as)\gamma v$ is a node of a global run of $\mathcal{M}'$, then $(p', as')uv$ and $(p_1, as'')\omega_1$ can be the right and left children of $(p, as)\gamma v$. For this, we add the transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \triangleright (p_1, as'')\omega_1$ into Δ'_i.

If the node $(p\omega, L)$ has the right and left children $(p'vu, L \cup \{l\})$ and $(p_2\omega_2, \emptyset)$ due to the transition rule $p\gamma \overset{acq(l)}{\hookrightarrow}_i p'v \rhd p_2\omega_2 \in \Delta_i$, then $l \notin L$ (see Item α_4). We suppose that the subtree rooted by $(p'vu, L)$ (resp. $(p_2\omega_2, \emptyset)$) in T be T_2 (resp. T_3) and the acquisition structure of T_2 (resp. T_3) is as' (resp. as''), then $l \in as'_X$, $as''_X = \emptyset$ and $as''_R = \emptyset$ according to the definition of the acquisition structures. Thus, if $l \notin as'_X \cup as''_X$, $Compatible(as', as'')$ does not hold, $as''_X \neq \emptyset$, $as''_R \not\emptyset$ or as' (resp. as'') is inconsistent, then, the global run T violates the lock usages. On the other hand, if $l \notin as'_R$ and $l \in as'_A \cup as''_A$, then the lock l acquired by applying $p\gamma \overset{acq(l)}{\hookrightarrow}_i p'v \rhd p_2\omega_2$ will finally acquired again before the release of l. This arises a deadlock. Thus, we do not add $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \rhd (p_1, as'')\omega_1$ into Δ'_i for any $as \in AS$ if ($l \notin as'_R$ and $l \in as'_A \cup as''_A$), or $l \notin as'_X \cup as''_X$, or $Compatible(as', as'')$ does not hold, or $as''_X \neq \emptyset$, or $as''_R \not\emptyset$ or as' (resp. as'') is inconsistent.

Otherwise, if $l \in as'_R$, $Compatible(as', as'')$, as' and as'' are consistent, according to the definition of the acquisition structures, the acquisition structure as of T_1 must be $\Big((as'_R \cup as''_R) \setminus \{l\}, ((as'_{RH} \cup as''_{RH}) \setminus (\mathbb{L} \times \{l\}) \cup (\{l\} \times (as'_R \cup as''_R) \setminus \{l\})), as'_U \cup as''_U \cup \{l\}, as'_{AH} \cup as''_{AH}, as'_A \cup as''_A, (as'_X \cup as''_X) \setminus \{l\}\Big)$. Indeed, the initial release of l together with $acq(l)$ will be a usage of l for T_1 instead of the initial release of l which implies that as_{AH} should not contain any edge $\mathbb{L} \times \{l\}$. This usage of l occurs after the other initial releases $as'_R \setminus \{l\}$. Thus, if $(p, as)\gamma v$ is a node of a global run of $\mathcal{M}'$, then $(p', as')uv$ and $(p_1, as'')\omega_1$ can be the right and left children of $(p, as)\gamma v$. For this, we add the transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \rhd (p_1, as'')\omega_1$ into Δ'_i.

Otherwise, if $l \notin as'_R$, $l \in as'_A \cup as''_A$, $Compatible(as', as'')$, as' and as'' are consistent, then, according to the definition of the acquisition structures, the acquisition structure as of T_1 must be $\Big(as'_R \cup as''_R, as'_{RH} \cup as''_{RH}, as'_U \cup as''_U, \{\{l\} \times (as'_U \cup as''_U)\} \cup as'_{AH} \cup as''_{AH}, as'_A \cup as''_A \cup \{l\}, (as'_X \cup as''_X) \setminus \{l\}\Big)$. Indeed, $acq(l)$ is a final acquisition of l due to $l \notin as'_R$ and the usages $as'_U \cup as''_U$ occur after this final acquisition of l. Thus, if $(p, as)\gamma v$ is a node of a global run of $\mathcal{M}'$, then $(p', as')uv$ and $(p_1, as'')\omega_1$ can be the right and left children of $(p, as)\gamma v$. For this, we add the transition rule $(p, as)\gamma \overset{\tau}{\hookrightarrow}_i (p', as')\omega \rhd (p_1, as'')\omega_1$ into Δ'_i.

Thus, we can get that $\mathcal{M}$ has a global run T starting from a global configuration $(p\omega, L)$ using locks in a nested style iff $\mathcal{M}'$ has a corresponding global run T' starting from a global configuration $(p, as)\omega$ for some $as \in AS$. By extending the valuation $\lambda : AP \longrightarrow 2^{\bigcup_{i=1}^{n}(P_i \times \Gamma_i^* \times 2^{\mathbb{L}})}$ to $\lambda' : AP \longrightarrow 2^{\bigcup_{i=1}^{n}(P_i \times AS \times \Gamma_i^*)}$, we can get the following theorem. Indeed, for every $ap \in AP$, $\lambda'(ap) = \{(p, as)\omega \in \bigcup_{i=1}^{n}(P_i \times AS \times \Gamma_i^*) \mid (p\omega, as_X) \in \lambda(ap)\}$.

Theorem 4. *Given a L-DPN $\mathcal{M} = (Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n$ s.t. for every i, $1 \leq i \leq n$, $\mathcal{P}_i = (P_i, \Gamma_i, \Delta_i)$, a LTL formula $f = \bigwedge_{i=1}^{n} f_i$ and a valuation λ, let $\mathcal{M}' = (\mathcal{P}'_1, ..., \mathcal{P}'_n)$ be the DPN such that for every $i : 1 \leq i \leq n$, the DPDS $\mathcal{P}'_i$ is computed as above. For every global configuration $(p_0\omega_0, \emptyset)$, $\mathcal{M}$ has a global run T starting from $(p_0\omega_0, \emptyset)$ that uses locks in a nested style and satisfies f iff $\mathcal{M}'$ has a global run T' starting from $(p_0, as)\omega_0$ for some $as \in AS$ such that T' satisfies f. Moreover, for every $i : 1 \leq$*

$i \leq n$, $|P'_i| = \mathbf{O}(|P_i| \cdot 2^{\mathbf{O}(|\mathbb{L}|^2)})$, $|\Delta'_i| = \mathbf{O}(|\Delta_i| \cdot 2^{\mathbf{O}(|\mathbb{L}|^2)})$, *and* $\mathcal{M}'$ *can be computed in time* $\mathbf{O}(\sum_{i=1}^{n} |\Delta_i| \cdot 2^{\mathbf{O}(|\mathbb{L}|^2)})$.

The complexity follows from the fact that the number of acquisition graphs AH and release graphs RH is at most $2^{\mathbf{O}(|\mathbb{L}|^2)}$ which implies that the number acquisition structures is at most $2^{\mathbf{O}(|\mathbb{L}|^2)}$.

By Theorem 4, Theorem 2 and Theorem 3, we can get the following two theorems. Note that the number of DCLICs in $\mathcal{M}'$ is at most $|\mathcal{D}_l| \cdot 2^{O(|\mathbb{L}|^2)}$.

Theorem 5. *Given a L-DPN* $\mathcal{M} = (Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n$*, a single-indexed LTL formula* $f = \bigwedge_{i=1}^{n} f_i$ *and a simple function* λ*, we can compute MAs* $\mathcal{A}_1, ..., \mathcal{A}_n$ *in time* $O(\sum_{i=1}^{n}(|\Delta_i| \cdot 2^{|f_i|} \cdot |\Gamma_i| \cdot |P_i|^3) \cdot 2^{O(|\mathbb{L}|^2)} \cdot 2^{|\mathcal{D}_l|} \cdot 2^{O(|\mathbb{L}|^2)})$ *s.t. for every* $i : 1 \leq i \leq n$ *and every* $(p\omega, L) \in P_i \times \Gamma_i^* \times 2^{\mathbb{L}}$*,* $(p\omega, L)$ *satisfies* f *iff there exist* $D \subseteq \mathcal{D}_{fp}$ *and* $as \in AS$ *s.t.* $((p, as)\omega, D) \in L(\mathcal{A}_{\pi(p)})$.

Theorem 6. *Given a L-DPN* $\mathcal{M} = (Act, \mathbb{L}, \mathcal{P}_1, ..., \mathcal{P}_n$*, a single-indexed LTL formula* $f = \bigwedge_{i=1}^{n} f_i$ *and a regular valuation* λ*, we can compute MAs* $\mathcal{A}_1, ..., \mathcal{A}_n$ *in time* $\mathbf{O}(\sum_{i=1}^{n}(|\Delta_i| \cdot 2^{|f_i|} \cdot |\Gamma_i| \cdot |States_i| \cdot |P_i|^3) \cdot 2^{O(|\mathbb{L}|^2)} \cdot 2^{|\mathcal{D}_l| \cdot 2^{O(|\mathbb{L}|^2)}})$ *s.t. for every* $i : 1 \leq i \leq n$ *and every* $(p\omega, L) \in P_i \times \Gamma_i^* \times 2^{\mathbb{L}}$*,* $(p\omega, L)$ *satisfies* f *iff there exist* $D \subseteq \mathcal{D}_{fp}$ *and* $as \in AS$ *s.t.* $((p, as)\omega, D) \in L(\mathcal{A}_{\pi(p)})$ *and* $as_X = L$*, where* $|States_i|$ *denotes the number of states of the automata corresponding to the regular valuation* λ.

4 Related work

DPNs and L-DPNs: The DPN model was introduced in [3]. Several other works use DPN and its extensions to model multi-threaded programs [3,8,18,20,25]. All these works only consider reachability issues. Ground Tree Rewrite Systems [9] and process rewrite systems [4,21] are two models of multi-threaded programs with procedure calls and threads creation. However, [21] only considers reachability problem and [9,4] only consider subclasses of LTL. We consider single-indexed LTL model checking problems.

Pushdown networks with communication between processes are studied in [2,6,1,23,5]. These works consider systems with a fixed number of threads. [16,17] use parallel flow graphs to model multi-threaded programs. However, all these works only consider reachability. [26] considers safety properties of multi-threaded programs.

Lock usages and acquisition structures: [14] first introduces (forward) acquisition histories that contains only the set of held locks and the acquisition graphs to check pairwise reachability properties of two threads (i.e. pushdown systems) communicating via well-nested locks. [12,13] extended the results of [14] with backwards acquisition histories that contains only the set of held locks and the release graphs to check fragments of LTL and CTL properties for two threads communicating via well-nested locks. [17] extended the acquisition histories of [14] to check pairwise reachability properties of programs with reentrant monitors (a restricted form of well-nested locks) and dynamic thread creation. [15] uses tuples-of-lock histories to check pushdown networks without threads creation for properties represented by a kind of finite automata. Using tuples-of-lock histories allows the decision procedure to use only one reachability query

of each pushdown system. While, in the worst case, [14] has to perform an exponential number of individual reachability queries of each pushdown system to handle an temporal operator. In order to compute predecessor sets of regular sets of configurations of L-DPNs, [19] introduces acquisition structures for L-DPNs. Their acquisition structures are similar to tuples-of-lock histories of [15] which are defined for pushdown networks without threads creation. [19] reduces the predecessor sets computation of L-DPNs to compute the predecessor sets of of DPNs [3].

Following [19], in this work, we reduce single-indexed LTL model-checking for L-DPNs to single-indexed LTL model-checking for DPNs. This latter problem can be solved by our previous work [22]. In this work, the acquisition structures we used are similar to the acquisition structures of [19]. However, in [19], the acquisition structures do not contain the set of held locks (i.e., as_X). and [19] first encodes the acquisition structures into a hedge automaton, a kind of finite automaton, and then computes the product DPN of the L-DPN with the hedge automaton, where the set of held locks are stored into the control locations of the DPN. Our work do not perform this intermediate procedure, i.e., encoding the acquisition structures into a hedge automaton. Instead, we directly compute the product DPN from a L-DPN and the acquisition structures. Our approach can remove inconsistent acquisition structures during this product. While, [19] removes them when computing the product DPN of the L-DPN and the acquisition structures. Moreover, during encoding of the acquisition structures into a hedge automaton, [19] checks only whether the sets of finally acquired locks of two acquisition structures are disjoint. when composing two threads. This may introduces some inconsistent acquisition structures (e.g., having cycles in a release graph) into the hedge automaton. While, in this work, we check the compatibility of two acquisition structures which can remove these inconsistent acquisition structures. This makes the DPN more small.

Recently, [10] introduces bounded lock chains, a generalization of well-nested locks and shows that pairwise reachability is decidable for pushdown networks (without threads creation) with bounded lock chains. [11] extends the results of [10] to show the decidability of the fragment LTL for pushdown networks without threads creation. [5] introduces contextual locking, another extension of well-nested locks, and shows that pairwise reachability for pushdown networks (without threads creation) with contextual locking. These works do not consider threads creation. It is unknown whether our approach checks single-indexed LTL properties for DPNs with lock chains or contextual locking. We leave them as future work.

References

1. M. F. Atig, A. Bouajjani, and T. Touili. On the reachability analysis of acyclic networks of pushdown systems. In *CONCUR*, pages 356–371, 2008.
2. A. Bouajjani, J. Esparza, and T. Touili. A generic approach to the static analysis of concurrent programs with procedures. In *POPL'03*. ACM, 2003.
3. A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *CONCUR*, pages 473–487, 2005.
4. L. Bozzelli, M. Kretínský, V. Rehák, and J. Strejcek. On decidability of ltl model checking for process rewrite systems. *Acta Inf.*, 46(1), 2009.

5. R. Chadha, P. Madhusudan, and M. Viswanathan. Reachability under contextual locking. In *TACAS*, pages 437–450, 2012.
6. S. Chaki, E. M. Clarke, N. Kidd, T. W. Reps, and T. Touili. Verifying concurrent message-passing c programs with recursive calls. In *TACAS*, pages 334–349, 2006.
7. J. Esparza, A. Kucera, and S. Schwoon. Model checking LTL with regular valuations for pushdown systems. *Inf. Comput.*, 186(2):355–376, 2003.
8. T. M. Gawlitza, P. Lammich, M. Müller-Olm, H. Seidl, and A. Wenner. Join-lock-sensitive forward reachability analysis for concurrent programs with dynamic process creation. In *VMCAI*, pages 199–213, 2011.
9. S. Göller and A. W. Lin. The complexity of verifying ground tree rewrite systems. In *LICS*, pages 279–288, 2011.
10. V. Kahlon. Boundedness vs. unboundedness of lock chains: Characterizing decidability of pairwise cfl-reachability for threads communicating via locks. In *LICS*, 2009.
11. V. Kahlon. Reasoning about threads with bounded lock chains. In *CONCUR*, pages 450–465, 2011.
12. V. Kahlon and A. Gupta. An automata-theoretic approach for model checking threads for ltl properties. In *LICS*, pages 101–110, 2006.
13. V. Kahlon and A. Gupta. On the analysis of interacting pushdown systems. In *POPL*, pages 303–314, 2007.
14. V. Kahlon, F. Ivancic, and A. Gupta. Reasoning about threads communicating via locks. In *Computer Aided Verification*, 2005.
15. N. Kidd, P. Lammich, T. Touili, and T. W. Reps. A decision procedure for detecting atomicity violations for communicating processes with locks. In *SPIN*, pages 125–142, 2009.
16. P. Lammich and M. Müller-Olm. Precise fixpoint-based analysis of programs with thread-creation and procedures. In *CONCUR*, pages 287–302, 2007.
17. P. Lammich and M. Müller-Olm. Conflict analysis of programs with procedures, dynamic thread creation, and monitors. In *SAS*, pages 205–220, 2008.
18. P. Lammich, M. Müller-Olm, and A. Wenner. Predecessor sets of dynamic pushdown networks with tree-regular constraints. In *CAV*, pages 525–539, 2009.
19. P. Lammich, M. Müller-Olm, and A. Wenner. Predecessor sets of dynamic pushdown networks with tree-regular constraints. In *CAV*, pages 525–539, 2009.
20. D. Lugiez. Forward analysis of dynamic network of pushdown systems is easier without order. *Int. J. Found. Comput. Sci.*, 22(4):843–862, 2011.
21. R. Mayr. Process rewrite systems. *Inf. Comput.*, 156(1-2):264–286, 2000.
22. F. Song and T. Touili. Model-checking dynamic pushdown networks. Technical report, `http://www.liafa.univ-paris-diderot.fr/~song/DPN.pdf`, 2013.
23. T. Touili and M. F. Atig. Verifying parallel programs with dynamic communication structures. *Theor. Comput. Sci.*, 411(38-39):3460–3468, 2010.
24. M. Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *J. Comput. Syst. Sci.*, 32(2):183–221, 1986.
25. A. Wenner. Weighted dynamic pushdown networks. In *ESOP*, pages 590–609, 2010.
26. E. Yahav. Verifying safety properties of concurrent java programs using 3-valued logic. In *POPL*, pages 27–40, 2001.